

CLAIMS

- 1 1. A method for enabling parity declustering in a balanced parity array of a storage sys-
2 tem, the method comprising the steps of:
 - 3 combining a plurality of unbalanced stripe arrays to form the balanced array, each
4 unbalanced stripe array having parity blocks on a set of storage devices that are disjoint
5 from a set of storage devices storing data blocks; and
 - 6 distributing assignment of storage devices to parity groups throughout the bal-
7 anced array.
- 1 2. The method of Claim 1 further comprising the step of, after a single or double storage
2 device failure, ensuring that all surviving data storage devices are loaded uniformly dur-
3 ing reconstruction of the failed storage device or devices.
- 1 3. The method of Claim 1 wherein the storage system is a filer.
- 1 4. The method of Claim 1 further comprising the steps of:
 - 2 dividing each storage device into blocks; and
 - 3 organizing the blocks into stripes across the devices, wherein each stripe contains
4 data and parity blocks from each of the devices of the balanced array.
- 1 5. The method of Claim 4 wherein the step of distributing comprises the step of select-
2 ing patterns of characters representing data storage devices of a stripe to thereby change
3 the association of the data storage devices with parity groups from stripe to stripe of the
4 balanced array.
- 1 6. The method of Claim 5 wherein the characters are binary numbers.
- 1 7. The method of Claim 5 wherein the characters are ternary numbers.

- 1 8. The method of Claim 1 further comprising the steps of:
- 2 configuring the balanced array as a RAID-4 style array;
- 3 initially under-populating the array with storage devices; and
- 4 adding storage devices until a fully populated array of predetermined size is
- 5 achieved.
- 1 9. The method of Claim 8 wherein the storage devices are disks.
- 1 10. A system that enables parity declustering in a balanced parity array of a storage sys-
- 2 tem, the system comprising:
- 3 a plurality of storage devices, each storage device divided into blocks that are
- 4 further organized into stripes, wherein each stripe contains data and parity blocks from
- 5 each of the devices of the balanced array;
- 6 a storage operating system including a storage layer configured to implement a
- 7 parity assignment technique that distributes assignment of devices to parity groups
- 8 throughout the balanced array such that all storage devices contain the same amount of
- 9 data or parity information; and
- 10 a processing element configured to execute the operating system to thereby in-
- 11 voke storage access operations to and from the balanced array in accordance with the
- 12 concentrated parity technique.
- 1 11. The system of Claim 10 wherein the storage layer further combines a plurality of un-
- 2 balanced stripe arrays to form the balanced array, each unbalanced stripe array having
- 3 parity blocks on a set of storage devices that are disjoint from a set of storage devices
- 4 storing data blocks.
- 1 12. The system of Claim 11 wherein the storage devices are disks and wherein the storage
- 2 layer is a RAID layer.
- 1 13. The system of Claim 12 wherein the RAID layer is implemented in logic circuitry.

- 1 14. The system of Claim 10 wherein the storage system is a network-attached storage ap-
2 pliance.
- 1 15. The system of Claim 10 wherein the storage devices are one of video tape, optical,
2 DVD, magnetic tape and bubble memory devices.
- 1 16. The system of Claim 10 wherein the storage devices are media adapted to store in-
2 formation contained within the data and parity blocks.
- 1 17. Apparatus for enabling parity declustering in a balanced parity array of a storage
2 system, the apparatus comprising:
3 means for combining a plurality of unbalanced stripe arrays to form the balanced
4 array, each unbalanced stripe array having parity blocks on a set of storage devices that
5 are disjoint from a set of storage devices storing data blocks; and
6 means for distributing assignment of devices to parity groups throughout the bal-
7 anced array such that all storage devices contain the same amount of data or parity infor-
8 mation.
- 1 18. The apparatus of Claim 17 further comprising:
2 means for dividing each storage device into blocks; and
3 means for organizing the blocks into stripes across the devices, wherein each
4 stripe contains data and parity blocks from each of the devices of the balanced array.
- 1 19. The apparatus of Claim 18 wherein the means for distributing comprises means for
2 selecting patterns of characters representing data storage devices of a stripe to thereby
3 change the association of the data storage devices with parity groups from stripe to stripe
4 of the balanced array.

- 1 20. A computer readable medium containing executable program instructions for ena-
2 bling parity declustering in a balanced parity array of a storage system, the executable
3 program instructions comprising program instructions for:
4 combining a plurality of unbalanced stripe arrays to form the balanced array, each
5 unbalanced stripe array having parity blocks on a set of storage devices that are disjoint
6 from a set of storage devices storing data blocks; and
7 distributing assignment of devices to parity groups throughout the balanced array
8 such that all storage devices contain the same amount of data or parity information.
- 1 21. The computer readable medium of Claim 20 further comprising program instructions
2 for:
3 dividing each storage device into blocks; and
4 organizing the blocks into stripes across the devices, wherein each stripe contains
5 data and parity blocks from each of the devices of the balanced array.
- 1 22. The computer readable medium of Claim 21 wherein the program instructions for
2 distributing comprises program instructions for selecting patterns of characters repre-
3 senting data storage devices of a stripe to thereby change the association of the data stor-
4 age devices with parity groups from stripe to stripe of the balanced array.